TITLE OF THE INVENTION

PROCESS FOR IMPLEMENTING A SWITCHED FULL-DUPLEX ETHERNET TYPE COMMUNICATION NETWORK WITH REDUNDANCY

5

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present patent document is a continuation of U.S. application serial no. 10/287,959 filed on November 4, 2002, and claims priority on French patent application FR 01 14264 filed on November 5, 2001, the entire contents of each of which are hereby incorporated herein by reference.

10

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

[0002] This invention relates to a process for implementation of a redundant switched full-duplex Ethernet type communication network, particularly in avionics.

15

DESCRIPTION OF THE RELATED ART

[0003] The Ethernet network, which is the reference in the world of communication networks, can be used to send data in digital form by packets or "frames", where a frame is defined as being a set of data sent in a single step on the network.

20   [0004] In an Ethernet network, the data in each frame are not interpreted. The network carries the data without understanding their meaning. A frame is composed of two types of data, network data that are used to route the frame to its correct destination, and useful data which comprise the "useful load" in the frame.

[0005] An Ethernet network is composed of different equipment that is subscribed to the

25   network, and connected to each other through a communication means formed of active equipment called switches, which perform three functions:

> connect network subscribers in point to point mode through
> physical links, which are physical supports for messages to be

30   exchanged, for example twisted pair cables,

> route (switch) frames sent by source equipment to one or more
> destination equipment,

check the integrity and the format of the Ethernet frame.

[0006] FIG. 1 illustrates an Ethernet network composed of two switches 11 interconnected to each other and each connected to three items of subscriber equipment 12 in point to point mode.

[0007] Operation of such a network is simple. Each network subscriber can send frames in digital form at any time towards one or several other subscribers. When a switch receives the frames, the "network information" data are analyzed to determine the destination equipment. The frames are then switched towards this equipment.

[0008] In the "switched full-duplex Ethernet type network" expression:

> the "full-duplex" term means that the subscriber can send and receive frames at the same time on the same link,

> the "switched" term means that the frames are switched in switches on appropriate outputs.

[0009] For example, this network may be a 100 Mbits/s switched full duplex type network on twisted pair; the term "twisted pair" means that connections between the equipment and the switches are composed of two pairs of cables, each pair being twisted; the term 100 Mbits/s simply means the transmission or reception speed of frames on the network.

[0010] The Ethernet technology imposes:

> a minimum size and a maximum size on the frames,

> an identification of the source and/or the destination(s) in each frame,

> a CRC ("Cyclic Redundancy Check") that checks the integrity of the transported data.

[0011] At the present time, in the civil aeronautics field, data exchanges between the various onboard computers are based on the use of the ARINC 429 aeronautical standard.

[0012] However, the switched full-duplex Ethernet network is frequently used in industry. The emergence of new communication technologies shows that this type of network is an open and standard solution (IEEE standard 802.3) with a considerable potential for

development as a local network. But this type of solution does not provide any means of guaranteeing segregation and transfer performances (in terms of throughput, latency, etc.) necessary for avionics applications.

[0013] The purpose of this invention is to improve the availability of such a network, by providing means of protecting against the loss of a link or a switch, to enable its use in avionics.

[0014] In general, in a distributed control system, the communication system is made redundant so that each node in a set of nodes (equipment) can be controlled from one of these nodes, and so that these nodes can be connected so as to increase the load factor on the system and therefore its efficiency.

[0015] A European patent application EP-0 854 610 describes an Ethernet communication redundancy process between a set of nodes forming such a distributed control system. These nodes are connected to each other in duplex through communication lines in a first and a second Ethernet network, that are independent from each other. At the transmission end, a first node transmits identical data on communication lines in the first and the second networks, a data identifier being added to the transmitted data. A second node at the reception end determines which of the identical data received from the first node through communication lines in the first and the second networks arrived first and uses it as the reception data. The second data are then rejected if they are identical to the first data.

[0016] Unlike the process described in this patent application which is applicable to data, the purpose of the invention is to make a process enabling frame by frame redundancy.

SUMMARY OF THE INVENTION

[0017] The invention relates to a process for implementing redundant switched full-duplex Ethernet type communication network including at least two independent elementary networks, each including at least one source subscriber equipment and at least one destination subscriber equipment, connected to each other through at least one physical link and through at least one switch, each equipment being connected to each elementary network. The process performs a frame by frame redundancy on each elementary network.

[0018] This process comprises in transmission:

adding a numbering field in each transmitted frame, to insert a frame number, and

3

sending this frame on each of the elementary networks.

[0019] It includes on reception:

storing the received frame number,

accepting this frame only if its number has not already been received.

[0020] Advantageously, the accepting of a frame takes place during a given time window.

[0021] In one advantageously embodiment, the virtual link concept is used, which is a conceptual view of a link from one source equipment to at least one destination equipment. A virtual link number is accepted in the numbering field of each transmitted frame. A virtual link is characterized by:

a transfer direction, the virtual link being single directional,

a source equipment,

one or several items of destination equipment,

a fixed passband,

a maximum guaranteed time for transfer of packets from a source equipment to a destination equipment,

a path fixed on the network,

a unique identifier.

[0022] Advantageously, the process according to the invention can be used for implementing a redundant switched full-duplex Ethernet type communication network in avionics.

[0023] This type of network redundancy, which may for example include doubling up the network, with each subscriber being connected to each of the two networks, one of the two packets being selected on reception, can increase the network availability; the network will continue to operate if it has one or several defective switches or links.

4

[0024] The invention can achieve redundancies of order 2 or more, independently of the communication stack and applications.

BRIEF DESCRIPTION OF THE DRAWINGS

[0025] A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when considered in connection with the accompanying drawings, wherein:

[0026] FIG. 1 illustrates an Ethernet network according to known art,

[0027] FIG. 2 illustrates the concept of a virtual link in an Ethernet network according to known art,

[0028] FIG. 3 illustrates an Ethernet network according to known art in which several virtual links are shown,

[0029] FIG. 4 illustrates services used in the process according to the invention,

[0030] FIG. 5 illustrates an example of second order redundancy in subscribed equipment according to the invention,

[0031] FIGS. 6A and 6B illustrate operation of the process according to the invention in transmission mode and in reception mode respectively.

DETAILED DESCRIPTION OF EMBODIMENTS

[0032] The invention relates to a process for the implementation of a redundant switched full-duplex Ethernet type communication network comprising at least 2 elementary networks each of which comprises at least one source subscriber equipment and at least one destination subscriber equipment connected to each other through at least one physical link. In the reminder of the description, we will consider two elementary networks RE1 and RE2 as a non-limitative example. This process makes a frame by frame redundancy on each of the elementary networks.

[0033] In transmission, it comprises:

adding a numbering field in each transmitted frame, to insert a frame number so that each frame can be identified in time,

sending this frame on each of the elementary networks.

[0034] In reception, it comprises:

5

storing the received frame number,

accepting this frame only if its number has not already been received.

[0035] Advantageously, the frame acceptance step only takes place during a given time window, so that only a limited memory size can be used, each frame number reappearing after a determined time.

[0036] Therefore, the process according to the invention only sends the first frame received from an elementary network, to the application considered, the other corresponding frames from other elementary networks being rejected.

[0037] FIG. 5 illustrates an example of a second order redundancy used in a subscriber equipment.

[0038] In one advantageous embodiment, the process according to the invention uses the virtual link concept to limit the end to end transfer time, in other words source equipment towards one or several destination equipment.

[0039] This virtual link (VL) concept provides means of isolating data transfers between a source equipment 13 and destination equipment 14. A virtual link VL is seen as a "pipe" on the network, as illustrated in FIG. 2.

[0040] A virtual link VL is characterized by:

a transfer direction, the virtual link being single directional,

a single source equipment 13,

one or several destination equipment 14,

a fixed passband (maximum number of packets and their size per second),

a maximum guaranteed time for transfer of packets from a source equipment 13 to a destination equipment 14, regardless of the behavior of the rest of the network, each virtual link having its own transfer time,

a path fixed on the network,

a unique identifier.

[0041] A network subscriber may comprise several virtual links VL1, VL2, VL3, as shown in FIG. 3. We have:

a virtual link VL1 from equipment 21 to equipment 23, 24 and 25,

5

a virtual link VL2 from equipment 21 to equipment 22 and 23,

a virtual link VL3 from equipment 23 to equipment 22.

[0042] When equipment 21 wants to send a packet to equipment 23, 24 and 25, it sends a packet on the virtual link VL1. When it wants to send a packet to equipment 22 and 23, it sends a packet on the virtual link VL2.

10   [0043] The difference between virtual links VL1 and VL2 is identified by the destination identifier in the packet. On the network, the virtual link to which a packet belongs is determined by the identifier of the virtual link in the packet.

[0044] A switch uses a static configuration table to determine the virtual links that it is required to switch, and the allowable number of packets for a virtual link.

15   [0045] The virtual link concept is a means of fixing communications between equipment by configuring routes and passbands allocated to the virtual links. Thus, the flow formed by a virtual link is sure to be not disturbed by other flows sharing the same physical links all along its route in the network.

[0046] Furthermore, the virtual link concept enables central flow management, to make sure

20   that the sum of the passbands allocated to virtual links on the same physical link does not exceed the capacities of the technology of this physical link. In the above example, the sum of the passbands of virtual links VL1 and VL2 must be less than the transmission capacity of the physical link from equipment 21.

[0047] Therefore, a virtual link is a conceptual representation of a link from a transmitter

25   equipment to one or several items of receiver equipment that have the following characteristics:

fixed and predetermined route on the network,

fixed and guaranteed passband,

maximum guaranteed end to end latency

7

explicit identification in a multidestination or "multicast" type Ethernet frame (multidestination MAC address) and this identification is kept when passing through one or more switches.

[0048] As illustrated in FIG. 4, the process according to the invention is then characterized by the implementation of several services in each subscriber equipment 50:

a transmission service, the role of which is to enable an application 52 to access virtual links in transmission (virtual links VL1 and VL2). This service multiplexes virtual links towards a physical link through an Ethernet interface, and for each virtual link sends packets as a function of the passband allocated to the virtual link.

A reception service 55 that decodes frames (virtual links VL21 and VL22), checks their format and makes useful data available to applications.

[0049] In these transmission and reception services, the application may treat a virtual link like a queue.

[0050] Other protection services help to guard against some network failures.

[0051] A service for protection of a passband in the switch, which for each incoming virtual link is capable of checking the time characteristics of packets (spacing between packets, the consumed passband). If the allowable characteristics are exceeded, the packets are simply destroyed to prevent a failure in a transmitter or a virtual link from adversely affecting traffic in other virtual links starting from this switch.

[0052] A subscriber network redundancy service 60 that sends and receives each frame on the two elementary networks RE1 and RE2, in order to implement network redundancy. This network duplication in two elementary networks RE1 and RE2, that is transparent for applications, provides a means of guarding against a failure in a switch or an interface (it does not replace system level redundancy). This network redundancy service 60 is connected to at least one first Ethernet interface 61 with the elementary network RE1, and a second Ethernet interface 62 with the elementary network RE2.

[0053] As shown in FIG. 5, the equipment in the process according to the invention to obtain redundancy of virtual links on the physical layer comprises at least two physical interfaces, so

that they can be connected to at least two independent elementary networks RE1 and RE2. Their communication stacks include redundancy mechanisms that enable:

> sending an identical frame on the elementary networks RE1 and RE2,

5

> selecting the first valid received frame.

[0054] Therefore, in this embodiment, the process according to the invention used in the subscriber equipment and applied by a virtual link onto the network are such that:

> in transmission, for each frame received from the communication stack:

10

> a numbering field is added so that a counter numbers the frame corresponding to each virtual link,

> this frame is sent onto the elementary networks RE1 and RE2.

> in reception, for each frame assigned to a virtual link:

> the frame number is stored,

15

> this frame is accepted if this number has not already been received, and if it is received it is destroyed.

[0055] FIG. 6A illustrates this embodiment of the process according to the invention in transmission, and 6B illustrates it in reception.

[0056] FIG. 6A illustrates the following in sequence:

20

> addition of a numbering field on the IP/UDP/data assembly,

> numbering of the frame for each virtual link,

> send the frame to each Ethernet controller that formats it as an IEEE 802.3 frame.

[0057] Numbering is done by virtual link, and the same number can be used for two different

25 virtual links.

[0058] FIG. 6B shows the following in sequence:

for each virtual link, verification of the number in the numbering field,

the frame is eliminated if the number has already been received,

otherwise the frame is sent to the higher layer.

5 [0059] The counter field can be small, and when the counter reaches its maximum value, frame numbering restarts from zero.

[0060] Obviously, numerous modifications and variations of the present invention are possible in light of the above teachings. It is therefore to be understood that within the scope of the appended claims, the present invention may be practiced otherwise than as specifically

10 described herein.